

Express Mailing Label No. EV 329195261 US

PATENT APPLICATION
Docket No. SLA1194

UNITED STATES PATENT APPLICATION

of

JON M. SPEIGLE

AND

JOHN E. DOLAN

for

**SYSTEMS AND METHODS FOR CORRECTING IMAGE COLOR
BALANCE**

SYSTEMS AND METHODS FOR CORRECTING IMAGE COLOR BALANCE

TECHNICAL FIELD

[01] The present invention relates generally to digital image processing and more particularly to methods and systems for selectively applying image illuminant corrections.

BACKGROUND

[02] Colors viewed in an image are dependent on the light that illuminates the subject of the image. Different illuminants will cause differences in the light reflected from the surfaces of the image subject matter and therefore differences in visual response to the reflected light. The human visual system approximately corrects these response differences to a perceived norm. However, when images are captured on media and viewed under a light source different than the source in the image scene, these natural corrections do not take place. Accordingly, recorded images need to be color-balanced to a "normal" light source in order to appear as they would to the natural eye. This balancing or color correction can be performed once the scene illuminant is identified.

[03] There are many known methods for identification of a light source or illuminant in an image scene. However, the conventional correction algorithms assume that all image pixels represent reflecting surfaces. When an image contains self-luminous objects such as sky and other light sources the surface-pixel assumption is violated. When an image contains a significant portion of non-reflective, self-luminous objects, conventional methods will fail and the image illuminant will be incorrectly determined. For example, if an image contains blue sky and the color-balance algorithm assumes that all pixels are reflecting objects, "bluish" pixels could be taken as evidence that the illumination of the scene is bluish. Because a color correction is approximately the opposite hue of the estimated illuminant, the correction for a bluish illuminant would be to shift the image in a yellowish direction. This correction might produce an overly yellowish ground/surface region and a desaturated sky region.

[04] These color correction, color balance or color constancy algorithms generally do not address the question of how to handle images containing luminous objects, which are also referred to herein as self-luminous objects. They have, rather, focused on images where the surface-pixel assumption is satisfied (e.g., uniformly illuminated Mondrian-like images).

[05] It would be advantageous if a method existed that permitted accurate digital image illuminant corrections to be made for images containing significant exception regions, such as self-luminous regions.

SUMMARY

[06] The present invention describes a method for addressing problems in application of color balance corrections to images containing exception characteristics, such as self-luminous or non-reflecting objects, color distribution properties and other parameters.

[07] Accordingly, a method of compensating digital image corrections is provided. The method comprises the acts of: calculating digital image correction parameters; determining image exception characteristics; and applying the corrections to the digital image in response to the image exception characteristics. The image exception characteristics can be factors such as detecting image self-luminous regions, color distribution properties, multiple illuminants and other parameters.

[08] Applying the corrections to the digital image in response to the image exception characteristics may comprise processes such as varying the attenuation of a correction, global omission of a correction, using a plurality of corrections, selective application of one or more corrections to spatial regions, selective application of one or more corrections to chromatic ranges, selective application of one or more corrections to luminance ranges, selective application of one or more corrections to self-luminous and/or reflective regions.

[09] For example, calculating digital image correction parameters may include calculating a correction for a set of pixels, and determining image exception characteristics may include detecting image self-luminous regions. Then, applying the corrections to the digital image in response to the image exception characteristics includes varying the attenuation of the correction for the set of pixels in response to pixel position or region. That is, a full, or near-full correction

is applied to pixel positions in non-luminous regions and a full, or near-full attenuation is applied to corrections in luminous regions.

[10] In some embodiments of the present invention, detecting image self-luminous regions includes detecting self-luminous regions in response to a combination of confidence measures such as pixel color characteristics, luminance characteristics and position or location characteristics. Other confidence measures may include combinations of the above-mentioned factors.

[11] For example, in some embodiments, a luminous region may be detected in response to bright pixel luminance, where bright pixel luminance is defined as either relatively bright with respect to surrounding image regions or brighter than a threshold brightness.

[12] In some embodiments, luminous regions may be detected in response to a function of projections of color values such as pixel chromaticity or the chromatic dimensions of a uniform colorspace such as CIE Lab. For example, a function based on a hue angle and/or distance from a reference point in a chromaticity chart may be used. In more specific embodiments, a CIE standard illuminant, such as Daylight 6500 (D65) can be used as a reference and a specific angle or angular range can be used to identify self-luminous regions or other regions of interest. In other embodiments color distance measures may also be used to identify or partially identify luminous characteristics.

BRIEF DESCRIPTION OF THE DRAWINGS

[13] The present embodiments will become more fully apparent from the following description and appended claims, taken in conjunction with the accompanying drawings. Understanding that these drawings depict only typical embodiments and are, therefore, not to be considered limiting of the invention's scope, the embodiments will be described with additional specificity and detail through use of the accompanying drawings in which:

Fig. 1a is a flowchart illustrating methods of self-luminous element detection of some embodiments of the present invention;

Fig. 1b is a flowchart illustrating methods of self-luminous element detection and element weighting of some embodiments of the present invention;

Fig. 2 is a an exemplary image used to illustrate elements of embodiments of the present invention;

Fig. 3a is an image illustrating a likelihood of being self-luminous based on pixel position relative to an image boundary;

Fig. 3b is an image illustrating a likelihood of being self-luminous based on color characteristics;

Fig. 3c is an image illustrating a likelihood of being self-luminous based on luminance characteristics;

Fig. 3d is an image illustrating a combined likelihood of being self-luminous based on the combination of factors illustrated in figures 3a, 3b and 3c;

Fig. 4 is a diagram illustrating an exemplary self-luminous probability distribution

Fig. 5 is a flowchart illustrating some steps of embodiments of the present invention;

Fig. 6 is a diagram illustrating the various color distributions reflected from selected illuminants;

Fig. 7 is a diagram illustrating the chromaticities of the illuminants shown in Fig. 6;

Fig. 8 is a flowchart illustrating methods of embodiments employing a weighting factor related to a likelihood of being self-luminous;

Fig. 9 is a flowchart illustrating methods of embodiments employing a correction factor;

Fig. 10 is a flowchart illustrating methods of embodiments wherein self-luminous elements are distinguished during correction;

Fig. 11 is a flowchart illustrating embodiments wherein image exception characteristics are used to modify correction parameters;

Fig. 12 is a flowchart illustrating embodiments wherein image exception characteristics comprising self-luminous characteristics are used to modify correction parameters;

Fig. 13 is a flowchart illustrating embodiments wherein image exception characteristics comprising image color gamut distribution characteristics are used to modify correction parameters;

Fig. 14 is a flowchart illustrating embodiments wherein image exception characteristics comprising multiple illuminant characteristics are used to modify correction parameters; and

Fig. 15 is a chromaticity diagram

DETAILED DESCRIPTION

[14] In some embodiments of the present invention, methods for determining image self-luminous regions are employed. Digital image correction parameters may be calculated in response to determining the self-luminous regions. Image self-luminous regions can be determined from image feature measures. For example, image feature measures can be pixel color characteristics, pixel luminance and/or pixel position parameters.

[15] In many algorithms of embodiments of the present invention, a color-balance correction takes the form of a matrix that is applied in a linear colorspace such as the tristimulus values of CIE 1931 XYZ. Input tristimulus values can be multiplied by a matrix to produce a set of color-corrected values as follows:

$$\mathbf{x}' = \mathbf{M} \cdot \mathbf{x}$$

where

\mathbf{x} is a vector of tristimulus values

\mathbf{M} is a color correction matrix

\mathbf{x}' is the vector of corrected values.

[16] A color adaptation model is often referenced in a color-correction context. In this model adaptation occurs independently within long-, medium-, and short-wavelength sensitive channels:

$$L' = g_L \cdot L$$

$$M' = g_M \cdot M$$

$$S' = g_S \cdot S$$

where L, M and S refer to the long-, middle-, and short-wavelength photoreceptor responses. In a cone colorspace the corrections could be implemented as a diagonal matrix. This same correction can be transformed by a change of basis to operate on CIE tristimulus values, in which

case it would still be 3-by-3 but would contain off-diagonal terms. Others have applied similar diagonal models for color correction in an RGB or the XYZ colorspace.

[17] A different method of defining a correction matrix is in terms of an estimated illuminant for the scene. These methods use linear models for reflectances and illuminants. After principle components analysis (or singular value decomposition, SVD) on a large set of reflectance or illuminant functions (e.g., those of the matte Munsell chipset or measured daylight power spectra), typical illuminants and surfaces can be described within low-dimensional linear models. Reflectance functions and illuminant spectral power distributions can then be approximated within the linear model by sets of weights:

$$\mathbf{r} = \mathbf{B}_s \mathbf{w}_s$$
$$\mathbf{e} = \mathbf{B}_e \mathbf{w}_e$$

where

\mathbf{r} is a 1-by-nWavelengths vector of reflectance values;

\mathbf{e} is a 1-by-nWavelengths vector of illuminant spectral power values;

\mathbf{w}_s is a set of weights within the surface linear model;

\mathbf{w}_e is a set of weights within the illuminant linear model;

\mathbf{B}_s is a set of surface basis functions; and,

\mathbf{B}_e is a set of illuminant basis functions.

[18] Sensor responses to a surface rendered under an illuminant can be computed as

$$\mathbf{x} = \mathbf{T} \mathbf{E} \mathbf{B}_s \mathbf{w}_s ,$$

where

\mathbf{E} is a diagonal [nWavelengths, nWavelengths] matrix representing the illuminant,
i.e., $\mathbf{E} = \text{diag}(\mathbf{B}_e \mathbf{w}_e)$;

\mathbf{T} is a matrix representing sensor spectral sensitivities, [nSensors-by-nWavelengths]

[19] The mapping from reflectance weights to sensor responses may be inverted using linear regression or some other error minimization technique:

$$\mathbf{w}_s = (\mathbf{T} \mathbf{E} \mathbf{B}_s)^{-1} \mathbf{x}$$

[20] and the reflectance weights may then be used to render the reflectance functions under a second reference illuminant. The equation relating original to corrected sensor responses is:

$$\mathbf{x}' = \mathbf{M} \cdot \mathbf{x}$$

where

$$\mathbf{M} = \mathbf{TE}_2 \mathbf{B}_s (\mathbf{TE}_1 \mathbf{B}_s)^{-1};$$

$$\mathbf{E}_2 = \text{diag}(\mathbf{e}_{\text{canonical}});$$

$$\mathbf{E}_1 = \text{diag}(\mathbf{B}_E \mathbf{w}_E);$$

$\mathbf{e}_{\text{canonical}}$ is the spectral power distribution of a canonical illuminant;

\mathbf{w}_E is an illuminant estimate for the original scene.

[21] In embodiments of the present invention, image elements, such as pixels or groups thereof, are evaluated to determine whether they are likely to violate assumptions of the colorbalance estimation method or assumptions of the colorbalance correction method. In the methods disclosed herein, colorbalance estimation and correction methods are made responsive to violations of the colorbalance algorithm assumptions.

[22] In some embodiments of the present invention, image elements, such as pixels or groups thereof, are evaluated to determine whether they are likely to represent self-luminous objects in the image. Such objects violate the assumption implicit in most colorbalance estimation and correction algorithms that image regions represent reflecting surfaces. Elements may be more likely to be self-luminous when they are located in proximity to the edges of the image. For example, and not by way of limitation, in images that contain regions of sky, a self-luminous area, the self-luminous sky area is typically located along the top edge of the image. When images are rotated, such as from landscape to portrait format, top edges may become side edges and vice versa. Accordingly, in some embodiments of the present invention, proximity to any edge can be an indicator of self-luminance. In other embodiments, image orientation algorithms may be used to identify the top, bottom and sides of an image. In other embodiments, an image may comprise metadata or some other indication of image orientation. In these embodiments with identified orientation, proximity to particular edges may be indicative of self-luminance. For example and not by way of limitation, a pixel's proximity to the top of the image may be more indicative of self-luminance than proximity to the side or bottom of the image.

[23] In some embodiments, an element's likelihood of being self-luminous may be determined in relation to the element color characteristics or descriptors. In some embodiments, the output

of this function can be viewed as a continuous-tone grayscale image, where lighter regions correspond to regions likely to be self-luminous objects and darker regions correspond to reflecting surfaces.

[24] In some embodiments, an element's chromaticity or non-intensive color characteristics can be used to determine self-luminance. When an element's chromaticity is similar to the chromaticity of known illuminants or self-luminous objects, then that element is more likely to be self-luminous.

[25] In some embodiments, an element's intensive characteristics can be used to determine self-luminance. When an element's intensive characteristic is similar to the intensive characteristic of known illuminants or self-luminous objects, then that element is more likely to be self-luminous.

[26] In some embodiments, an element's intensive and non-intensive characteristics can jointly be used to determine self-luminance.

[27] Some embodiments of the present invention comprise the steps illustrated in the chart of Fig.1a wherein an image element, such as a pixel or region of interest, is evaluated 10 to determine its characteristics such as location, non-intensive color characteristics such as chromaticity and/or intensive values such as luminance. The location of the element is evaluated 12 to determine its proximity to image edges and, if the element is sufficiently proximate to an edge, a specific edge or some other geometric image boundary the element is considered more likely to be self-luminous.

[28] The color characteristics of the element are then evaluated 14 and the element is considered more likely to be self-luminous if the color characteristics match color characteristics of known illuminants or self-luminous objects.

[29] The results of these evaluations can be combined 16 and, if sufficient evidence exists, the element can be considered likely to be self-luminous 18.

[30] In other embodiments of the present invention, as illustrated in Fig. 1b, an image element can be evaluated and classified according to a weighting factor or score. In these embodiments, an element is examined to determine whether its location is proximate to image edges, specific edges or some other image boundary characteristic. If the location is sufficiently proximate to

one of more of these characteristics, the weight factor or score of that element is increased 24 and the element is further evaluated.

[31] If an element's color characteristics are similar to those of known self-luminous objects 26, such as blue sky, overcast sky, tungsten-incandescent light, fluorescent light and others, that element can be considered likely to be self-luminous and the element's weighting factor or score can be increased 28 to reflect this likelihood.

[32] If an element's luminance characteristics are found to be similar to those of known self-luminous objects 30, the weighting factor or score of that element can be increased 32 to reflect a higher likelihood that the element is self-luminous. When the score or weighting factor of an element is sufficiently high, the element may be considered self-luminous.

[33] The methods of some embodiments of the present invention may be illustrated with reference to Fig. 2, which is an original digital image and Figs. 3a-3d, which are a set of images derived from Fig. 2 depicting the likelihood of self-luminosity based on element location, chromaticity or color characteristics, and/or luminance parameters. In the images in Figures 3a-3d, white areas represent areas of high likelihood, while black areas represent low likelihood.

[34] In some embodiments of the present invention, a self-luminosity decision variable, s , can be computed as:

$$s = f(x, i, j)$$

where

x represents a neighborhood of color values; and

i, j represent the position of the neighborhood

[35] In some embodiments, independence between the factors can also be assumed. This is convenient computationally at the expense of not taking advantage of correlations between terms (e.g., that sky is likely to be both bluish and of high luminance).

$$s = g(x) * h(i, j)$$

[36] In some embodiments, the term dependent on color values can be decomposed into separable terms based on chromaticity, c , and luminance, Y :

$$s = g_1(c) * g_2(Y) * h(i,j)$$

[37] In some embodiments, $g_1(c)$ can be designed such that a maximum is reached at the hue angle of "blue sky" and that decreases as the hue angle deviates from this value. Similarly, $g_2(Y)$ can be based on the luminance distribution for a large number of sky regions. For $h(i,j)$, the simple assumption can be made that sky is more likely near the edges of the image. Here $h(i,j)$ is a linear function of pixel row number. Note that these choices for $g_1(c)$ and $g_2(Y)$ strongly bias the model toward identifying regions of "blue sky" as the self-luminous objects.

[38] Fig. 2 is an image containing surface regions 42 and sky regions 40. Fig. 3 illustrates maps of self-luminous likelihood functions. We use "xy" to denote chromaticity and "Y" to denote luminance, after the CIE 1931 xyY colorspace. In the ground regions, the final probability is near zero as illustrated by the black color in all four images in Figs 3a-3d.

[39] In Fig. 3a, the likelihood of an element being self-luminous near the top edge of the image is shown by the lighter shading that gets progressively darker as the location becomes more distant from the top edge.

[40] In Fig. 3b, the elements or pixels that have chrominance similar to "blue sky" are shown in white, while those "distant" from that range of hues are shown in darker shades.

[41] In Fig. 3c, the elements or pixels that have luminance values similar to those of known self-luminous objects are shown in lighter shades. In this exemplary image case, few pixels have high luminance as the sky in this image is relatively dark.

[42] Fig. 3d shows a composite image representing a combination of self-luminosity feature images shown in Figs. 3a-3c. In this figure, only elements in the sky are shown in lighter shades thereby signifying a successful application of the combined analysis.

[43] Fig. 4 is a diagram illustrating the probability distribution corresponding to Fig. 3d. The "sky" region corresponds to the values in the 0.2 - 0.35 range.

[44] Embodiments of the present invention may be used in conjunction with color balance or color constancy methods and correction of color balance in digital images. In some embodiments of the present invention, as shown in Fig. 5, inventive methods and systems may be used to determine self-luminous elements 110. These methods and systems may further influence the determination of an image illuminant 115 and the correction used to achieve color

balance 120. These systems and methods may also have an effect on the application of the color balance correction 130.

[45] Most color-balance algorithms can be modified to use a self-luminous classification map to derive color-balance parameter estimates from only those pixels classified as non-self-luminous. These modified algorithms would produce estimates with reduced bias due to the presence of self-luminous image regions.

[46] For example, the greyworld algorithm could be modified to compute a weighted average of the scene chromaticities where the weight would be reduced for pixels likely to be self-luminous.

[47] Consider the illuminant estimation method that selects the illuminant parameters e_i such that the function $f(e_i)$ is maximized for the image, I :

$$f(e_i) = \sum p(x_k | I, S_k=0) * p(x_k | e_i)$$

where

x_k represents a vector of color values

S_k represents a binary variable indicating the state of self-luminous ($S=1$) or surface ($S=0$)

e_i represents a set of parameters for the i th illuminant

$p(x_k | I, S_k=0)$ represents the likelihood of observing color coordinates x_k in image I for the non-self-luminous pixels.

$p(x_k | e_i)$ represents the likelihood of observing color coordinates x_k under illuminant e_i .

This term represents a scene/image color probability model that is indexed by illuminant.

[48] When the color frequencies observed in the image match the color frequencies expected under a particular illuminant, the value f will be large. The illuminant corresponding to the maximum of $f(e_i)$ is the illuminant most consistent with the surface pixels in the image.

[49] For a discrete set of color coordinates, the observed color likelihood function can be approximated as:

$$p(x_k | I, S=0) = n_k / N_{S=0}$$

where

n_k represents the number of surface pixels within a distance ϵ of coordinate x_k

$N_{S=0}$ is the number of pixels in the surface class

[50] An alternative method of estimating an illuminant from an image, I , is to directly use the self-luminous decision variable rather than a classification map. Consider the illuminant estimation method that selects the illuminant parameters e_i such that the function $g(e_i)$ is maximized:

$$g(e_i) = \sum w_k \cdot p(x_k | e_i)$$

where

w_k represents a weight for a k th vector of color values that is dependent on the corresponding self-luminous decision variable

[51] The value w_k may be computed from pixel self-luminous decision variable, s :

$$w_k = 1 - s$$

where s is normalized to have range 0 – 1.

[52] $p(x_k | e_i)$ represents the likelihood of observing color coordinates x_k under illuminant e_i . The term w_k weights the contributions of pixels that are non-self-luminous more than pixels that are self-luminous. Thus the value $g(e_i)$ will be a measure of how consistent illuminant e_i is with the surface pixels in the image.

[53] For a discrete set of color coordinates, the weight may be computed as:

$$w_k = \sum w_j$$

where

j indexes over the pixels within a distance ϵ of coordinate x_k

[54] In alternative embodiments, a Bayesian color balance estimation technique can be used to select the illuminant, e_i , that maximizes the illuminant posterior probability distribution, $p(e_i | I)$, of the illuminant parameters conditioned on the image data. Bayes' Law relates the posterior to the prior, $p(e_i)$, and image likelihood function, $p(I | e_i)$:

$$p(e_i | I) = k \cdot p(I | e_i) \cdot p(e_i)$$

where

k is a normalizing constant

$p(e_i)$ represents the likelihood of different illuminants

$p(I | e_i)$ represents the likelihood the image given that the scene was illuminated by illuminant e_i .

[55] If pixels are assumed independent from one another, then the color likelihood of the image is the product of the likelihoods of each pixel's color values, x_j , given the illuminant:

$$p(I | e_i) = \prod p(x_j | e_i)$$

[56] One method of introducing a self-luminous classification map into the pixel likelihood function is to let a binary variable S indicate that a pixel is self-luminous ($S=1$) or non-self-luminous ($S=0$). The color likelihood function is equal to,

$$p(x_j | e_i) = p(x_j | e_i, S=0)p(S=0) + p(x_j | e_i, S=1)p(S=1)$$

[57] For the case where a pixel is self-luminous, the color likelihood function can be treated as conditionally independent of the illuminant:

$$p(x_j | e_i, S=1) = p(x_j | S=1)$$

where

$p(x_j | S=1)$ is a probabilistic model for self-luminous regions

Substituting, the image likelihood is computable as:

$$p(I | e_i) = \prod \{p(x_j | e_i, S=0)p(S=0) + p(x_j | S=1)p(S=1)\}$$

From the image likelihood function, the illuminant posterior distribution may be computed and then maximized to yield an illuminant estimate.

[58] In some embodiments, the posterior, $p(e_i | I)$, may be computed by multiplying the image likelihood and illuminant prior. The prior, $p(e_i)$, may be used to account for the relative likelihood of different illuminants; for instance, illuminants may be expected to be near the daylight locus.

[59] In some embodiments, a cost function may be used and computed by multiplying the posterior by a function which modifies the result toward a reference illuminant (such as CIE D6500). A bias toward a reference illuminant can be useful because corrections associated with extreme estimated illuminants can produce undesirable resultant images if the illuminant estimate is in error.

[60] The illuminant estimate may be the illuminant associated with the maximum after applying the cost function.

[61] The color likelihood function given the illuminant parameters, $p(x_j | e_i)$, is computed by making assumptions or measurements of the distributions of colors in scenes. The purpose of this term is to characterize the likelihood of different colors under each illuminant. For example, under a reddish illuminant chromaticities with a nominally reddish hue would be more likely and nominally bluish chromaticities would be highly unlikely.

[62] Fig. 6 illustrates how a set of matte Munsell surfaces shifts under a set of illuminants that are near the daylight locus. Each column of plots corresponds to a different illuminant. There is a clear shift in the chromaticities of the Munsell set with each illuminant.

[63] Fig. 7 illustrates the chromaticities of the illuminants depicted in Fig. 6. The color of these illuminants varies in a bluish-yellowish direction.

[64] One method of computing $p(x_j | e_i)$ is to assume that colors in images are uniformly drawn (with replacement) from the set of Munsell matte patches. The long-term average color frequencies will be the same as rendering the Munsell set under the illuminant. The probability of a given chromaticity, $[x,y]$, under the i th illuminant will simply be the frequency at that chromaticity of the distribution of the Munsell set under that illuminant. An alternative method is to derive an empirical color model from the surface color distributions across a large collection of images or objects.

[65] In some embodiments of the present invention, a likelihood of occurrence of specific colors may be pre-computed for a set of illuminants. These may be referred to as model gamuts, $p(x | e_i)$. The model gamuts represent the probability of occurrence of each color given the illuminant.

[66] In some embodiments, the probability of being self-luminous may be computed for each pixel as a function of color and position, $p(S=1 | x, i, j)$. The inverse value, $p(S=0 | x, i, j)$, may be accumulated as a function of color value to form a reduced image self-luminosity feature vector, $q(x_k)$ specified at a set of color coordinates.

[67] Illuminant parameters e_i may then be selected such that the function $F(e_i)$ is maximized for the image, I :

$$F(e_i) = \sum q(x_k) * p(x_k | e_i)$$

[68] In some embodiments, the image likelihood, $p(I | e_i)$, may be approximated as the inner product between $q(x_k)$ and $p(x_k | e_i)$. From this approximation, an approximation to the illuminant posterior distribution can be computed which combines the approximated image likelihood with the illuminant prior. This approximated posterior distribution may then be modified by an illuminant cost function and subsequently maximized to select an estimated illuminant.

[69] Some embodiments of the present invention may be explained with reference to Fig. 8, which is a flowchart illustrating the steps of an image illuminant process. Although the method is depicted as a sequence of numbered steps for clarity, no order should be inferred from the numbering unless explicitly stated. It should be understood that some of these steps may be skipped, performed in parallel, or performed without the requirement of maintaining a strict order of sequence. Initially, image element characteristics are obtained 200. These characteristics typically comprise image element color characteristics, luminance characteristics and location data as well as other characteristics. Using the element characteristic data, the likelihood that each image element is self-luminous can be determined. This may be performed as explained in other embodiments above with reference to image element proximity to image edges or boundaries, to element chromaticity relative to chromaticity of known illuminants and to element luminance.

[70] A weighting factor may be assigned 204 to each element based on its likelihood of being self-luminous. The most likely illuminant is then estimated for each element based on element characteristics 206. The image illuminant is then estimated 208 using image element characteristics and the weighting factors. The image illuminant may be estimated 208 in many ways according to the methods of embodiments of the present invention. In some embodiments, the weighting factor may be used to adjust the contribution of elements to the image illuminant estimation process. This adjustment may comprise complete omission of likely self-luminous elements or diminution of their contribution to the estimation process.

[71] Other embodiments of the present invention may be explained with reference to Fig. 9, which is a flow chart showing an image color-balance correction process. In these embodiments,

image element characteristics are obtained 220 in a typical manner as is known in the art. A likelihood that each image element is self-luminous is then determined 222 according to methods explained herein. A global image illuminant is estimated 224 for the entire image based on image element characteristics and the likelihood that each element is self-luminous. A correction factor is then computed 226 to color-balance the image for the illuminant that has been estimated. This correction factor is then applied to the image 228 to achieve proper color-balance.

[72] Other embodiments of the present invention may be explained with reference to Fig. 10, which is a flow chart illustrating a method for selective application of a color-balance correction factor. In these embodiments, image element characteristics are obtained 240 by methods known in the art. A likelihood that each element is self-luminous is determined 242 as explained above and an image illuminant is estimated 244 based on the element characteristics and the likelihood that each element is self-luminous. One or more correction factors are computed 246 to correct the image for the illuminant and the correction factor or factors are applied to the image 248 according to each element's likelihood of being self-luminous. In these embodiments, the correction factor may be applied only to elements that are not likely to be self-luminous, may be applied in degrees according to magnitude of the likelihood of being self-luminous or may be applied in other ways that are proportional or otherwise related to the likelihood of an element being self-luminous.

[73] Further embodiments may comprise estimating a modified common illuminant. In these embodiments a map of reflective regions and self-luminous regions is created. A composite illuminant is calculated primarily in response to the reflective regions and secondarily in response to de-emphasized self-luminous regions. A correction for this composite illuminant is then computed that shifts the color values of the pixels in the reflective and self-luminous regions in response to the difference between the composite illuminant and the selected reference illuminant.

[74] In some embodiments, determining an illuminant of a reflective region comprises sorting color values into a set of groups and counting the number of members of each color value group. The illuminant most likely to generate the set of color value groups is then selected.

[75] In alternative embodiments, determining an illuminant of a reflective region includes sorting color values into a set of groups distinguished by pixel chromaticity and counting the number of members of each chromaticity group. The illuminant most likely to generate the set of chromaticity groups is then selected.

[76] Some embodiments of the present invention may be explained in reference to Fig. 11, which is a flowchart illustrating methods of modifying digital image corrections. Although the methods are depicted as a sequence of numbered steps for clarity, no order should be inferred from the numbering unless explicitly stated. It should be understood that some of these steps may be skipped, performed in parallel, or performed without the requirement of maintaining a strict order or sequence.

[77] In these embodiments, digital image correction parameters are calculated 300. These correction parameters may be related to color balance, illuminant detection, or other image enhancement, image processing or similar processes. Image exception characteristics are also determined 302. These exception characteristics may be related to image illuminants, image color characteristics, image luminance characteristics, image chrominance ranges such as small chrominance gamut or unlikely chrominance gamut and other image characteristics. When these exception characteristics are determined, the image correction parameters may be modified 304 in response to these exception characteristics.

[78] In some embodiments, modification in response to the exception characteristics may comprise differential application of the correction parameters to segments of the digital image, varying the attenuation of a correction, using a plurality of corrections and other modifications and processes.

[79] In some embodiments of the present invention, illustrated in Figure 12, calculating digital image correction parameters 300 comprises estimating an image illuminant 310 and calculating a color-balance correction 312. Determining image exception characteristics 302 may also comprise detecting self-luminous regions 314 and modifying corrections in response to exception characteristics 304 comprises modifying the color-balance correction for application in the self-luminous regions 316. In some of these embodiments, modifying the color-balance correction 316 may comprise attenuation of the color-balance correction in self-luminous regions, omission

of the correction in self-luminous regions, application of a different correction in self-luminous regions and other correction modification techniques.

[80] In some embodiments of the present invention, illustrated in Figure 13, calculating digital image correction parameters 300 comprises estimating an image illuminant 320 and calculating a color-balance correction 322. Determining image exception characteristics 302 may also comprise detecting image color gamut characteristics 324 such as a color gamut range or a color gamut likelihood of occurrence. In these embodiments, modifying corrections in response to exception characteristics 304 may comprise modifying the color-balance correction in relation to the color gamut characteristics 326. In some of these embodiments, modifying the color-balance correction 326 may comprise attenuation of the color-balance correction when a color gamut is unusually restricted or unlikely to occur in a specific type of image, omission of the correction when a color gamut is unusually restricted or unlikely to occur in a specific type of image, application of a different correction when a color gamut is unusually restricted or unlikely to occur in a specific type of image and/or other correction modification techniques.

[81] In some embodiments of the present invention, illustrated in Figures 11 and 14, calculating digital image correction parameters 300 comprises estimating an image illuminant 330 and calculating a color-balance correction 332. Determining image exception characteristics 302 may also comprise detecting image illuminant characteristics 334 and modifying corrections in response to exception characteristics 304 may comprise detection of multiple illuminants 336. When multiple illuminants are detected 336 a correction may be modified 340. When only a single illuminant is detected 336, the correction may be applied to the image without modification 338.

[82] In some of these embodiments, modifying the color-balance correction 340 may comprise attenuation of the color-balance correction in an entire image or in regions in which additional illuminants are detected, omission of the correction in an entire image or in regions where additional illuminants are detected, application of different corrections in regions illuminated by different illuminants and other correction modification techniques.

[83] In some embodiments, the attenuation of a correction may be varied in response to pixel position, this step may comprise applying the full or partially attenuated correction to pixel

positions in non-luminous regions and fully attenuating the correction in luminous regions. In other embodiments, varying the attenuation of a correction in response to pixel position comprises linearly attenuating the correction as the pixel position changes from a non-luminous region to a luminous region. Alternately, the correction is non-linearly attenuated as the pixel position changes from a non-luminous region to a luminous region.

[84] In some embodiments, a correction may be differentially applied to various regions or areas of an image. This differential application may comprise vetoing or omitting a correction globally for the entire image. Differential application may also comprise a region-based correction. In some embodiments, the correction may be strongly attenuated in self-luminous regions, small gamut regions, unlikely gamut regions, multiple illuminant regions or other regions, while a full or partially attenuated correction is applied in other regions of the same image. Differential application may also comprise a spatially-varying correction responsive to a map, where a different illuminant is corrected in each region of the map.

[85] In some embodiments, detecting image self-luminous regions 314 may comprise detecting self-luminous regions in response to a combination of confidence measures such as pixel color values and position parameters. For example, self-luminous regions may be detected in response to a function of chromaticity or color characteristics.

[86] Some embodiments of the present invention may be described in relation to Fig. 15, which is a diagram 350 of CIE x-y chromaticity space. Points within the diagram 350 represent specific chromaticity values. The chromaticity of various illuminants may be represented by points 352 on the diagram. Specific points may represent reference points or values, such as reference point 354, which represents Daylight 6500 (D65). The diagram 350 may be implemented as a 2D lookup table from pixel chromaticity to weight. One method of creating this lookup table uses the chromaticity distribution of self-luminous objects across a large set of images. Using mainly sky images for example, the chromaticity distribution can be approximated parametrically with an angular function 356 relative to a chromaticity reference point, such as reference point D65 (354). In these embodiments, the bluish chroma of the daytime sky within an area that can be represented by an angular range radiating from reference point D65 (354). In some embodiments, a chromaticity range may be represented by a distance

from a reference point on the chromaticity diagram 350 or by a combination of angle and distance, such as with a polar coordinate based on a reference point, such as D65 (354).

[87] Once the chromaticity or other color characteristics of known self-luminous objects are known, pixel characteristics can be analyzed to determine similarities with known self-luminous objects.

[88] In other embodiments, these methods can be based upon the frequencies of the raw self-luminous pixels. In other embodiments, an illuminant chromaticity may be modeled as a function of sun/camera relative angle and atmospheric condition. As the angle becomes closer to the sun, the illumination spectral power would shift from bluish-to-yellowish/whitish.

[89] In further embodiments of the present invention, self-luminous regions may be detected in response to the position of a pixel with respect to image edges. One of the most common self-luminous objects is the sky. In many digital photographs and other images, the sky is located at the top of the image adjacent to the top edge. When orientation is unknown, the sky may be adjacent to any edge. Accordingly, self-luminous objects, such as the sky, are typically located adjacent to image edges. This position-based likelihood can be used to improve self-luminous object detection.

[90] In still other embodiments, self-luminous regions may be detected in response to a confidence measure function such as pixel luminance. For example, self-luminous regions may be detected in response to a relatively high pixel luminance value, where high pixel luminance is defined as relatively high with respect to other image pixels, or higher than a threshold luminance.

[91] Similarly to the chromaticity-based term discussed above, some embodiments may use several ways of defining a luminance-based term. The empirical approach is to look at the luminance distribution for known self-luminous regions from a large set of images. Approximations to any empirical data like this may also be valid. This term can be implemented using an equation or using a 1D look-up table (LUT). Some embodiments of the present invention concatenate a 2D chromaticity mapping and a 1D luminance mapping. Other embodiments may use a full 3D mapping when this complexity is appropriate. A full 3D

function mapping from pixel color to probability of being self-luminous may be employed in some embodiments.

[92] In some embodiments, detecting self-luminous regions in response to a combination of confidence measures comprising pixel color values, pixel position parameters and pixel luminance characteristics may comprise differentially weighting a pixel self-luminance confidence measure based on these characteristics.

[93] What is claimed is: